

# **Exhibit E**

UNITED STATES PATENT AND TRADEMARK OFFICE  

---

BEFORE THE PATENT TRIAL AND APPEAL BOARD  

---

GOOGLE LLC,  
Petitioner,

v.

SINGULAR COMPUTING LLC,  
Patent Owner.

---

IPR2021-00179  
Patent 8,407,273 B2  

---

Before JUSTIN T. ARBES, KRISTI L. R. SAWERT, and  
JASON M. REPKO, *Administrative Patent Judges*.

PER CURIAM.

DECISION  
Granting Institution of *Inter Partes* Review  
35 U.S.C. § 314

I. INTRODUCTION

*A. Background and Summary*

Petitioner Google LLC filed a Petition (Paper 2, “Pet.”) requesting *inter partes* review of claims 1–26, 28, 32–61, 63, and 67–70 of U.S. Patent No. 8,407,273 B2 (Ex. 1001, “the ’273 patent”) pursuant to 35 U.S.C. § 311(a). Patent Owner Singular Computing LLC filed a Preliminary Response (Paper 9, “Prelim. Resp.”) pursuant to 35 U.S.C. § 313. Petitioner

IPR2021-00179

Patent 8,407,273 B2

also filed an explanation for filing multiple petitions ranking its petition in Case IPR2021-00178 ahead of its Petition in this proceeding (Paper 3).

With our authorization (Paper 13), Petitioner also filed a Reply (Paper 14, “Reply”) and Patent Owner filed a Sur-Reply (Paper 15, “Sur-Reply”).

Pursuant to 35 U.S.C. § 314(a), the Director may not authorize an *inter partes* review unless the information in the petition and preliminary response “shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition.” For the reasons that follow, we institute an *inter partes* review as to claims 1–26, 28, 32–61, 63, and 67–70 of the ’273 patent on all grounds of unpatentability asserted in the Petition.

### *B. Related Matters*

The parties indicate that the ’273 patent is the subject of *Singular Computing LLC v. Google LLC*, Case No. 1:19-cv-12551-FDS (D. Mass.). *See* Pet. xi; Paper 8, 1. Petitioner filed another petition challenging claims 1–70 of the ’273 patent in Case IPR2021-00178, and filed four other petitions challenging claims of two related patents also asserted in the district court case in Cases IPR2021-00154, IPR2021-00155, IPR2021-00164, and IPR2021-00165.

### *C. The ’273 Patent*

The ’273 patent, entitled “Processing with Compact Arithmetic Processing Element,” issued on March 26, 2013. Ex. 1001, codes (45), (54). The ’273 patent describes “computer processors or other devices which use low precision high dynamic range (LPHDR) processing elements to perform computations (such as arithmetic operations).” *Id.* at col. 5, l. 65–col. 6, l. 2.

IPR2021-00179

Patent 8,407,273 B2

According to the '273 patent, conventional central processing unit (CPU) chips make inefficient use of transistors as a tradeoff for delivering the high precision required by many applications. *Id.* at col. 3, ll. 7–22. For example, conventional CPU chips “perform[] exact arithmetic with integers typically 32 or 64 bits long and perform[] rather accurate and widely standardized arithmetic with 32 and 64 bit floating point numbers,” but require “on the order of a million transistors to implement the arithmetic operations.” *Id.* at col. 3, ll. 15–22. According to the '273 patent, “many economically important applications . . . are not especially sensitive to precision and . . . would greatly benefit, in the form of application performance per transistor, from the ability to draw upon a far greater fraction of the computing power inherent in those million transistors.” *Id.* at col. 3, ll. 23–28. But “[c]urrent architectures for general purpose computing fail to deliver this power.” *Id.* at col. 3, ll. 28–29.

The '273 patent is, therefore, “directed to a processor or other device, such as a programmable and/or massively parallel processor or other device, which includes processing elements designed to perform arithmetic operations . . . on numerical values of low precision but high dynamic range (‘LPHDR arithmetic’).” *Id.* at col. 2, ll. 11–18. According to the '273 patent, “‘low precision’ processing elements perform arithmetic operations which produce results that frequently differ from exact results by at least 0.1%.” *Id.* at col. 2, ll. 28–31. In addition, “high dynamic range” processing elements “are capable of operating on inputs and/or producing outputs spanning a range at least as large as from one millionth to one million.” *Id.* at col. 2, ll. 35–39.

IPR2021-00179

Patent 8,407,273 B2

Figure 6 of the '273 patent is reproduced below.

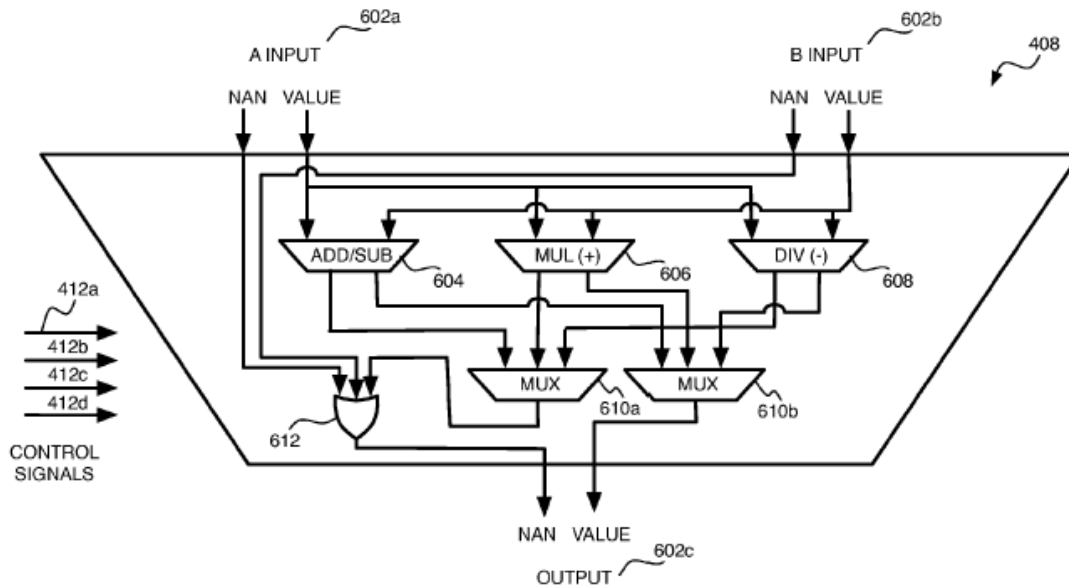


FIG. 6

Figure 6 depicts “an example design for an LPHDR arithmetic unit according to one embodiment of [the '273 patent].” *Id.* at col. 2, ll. 56–57. As shown in Figure 6, LPHDR arithmetic unit 408 receives two inputs: A input (602a) and B input (602b), and produces an output 602c. *Id.* at col. 12, ll. 52–53. The LPHDR arithmetic unit “is controlled by control signals 412a-d, coming from the CU 106, that determine which available arithmetic operation will be performed on the inputs 602a-b.” *Id.* at col. 12, ll. 59–62. According to the '273 patent, Figure 6 illustrates an embodiment where “all the available arithmetic operations are performed in parallel on the inputs 602a-b by adder/subtractor 604, multiplier 606, and divider 608.” *Id.* at col. 12, ll. 62–65. Finally, multiplexers (MUXes) 610a and 610b choose and send the desired result from among the outputs of the adder/subtractor, multiplier, and divider to output 602c. *Id.* at col. 13, ll. 1–8.

IPR2021-00179

Patent 8,407,273 B2

The '273 patent provides that “[t]he computing architecture literature discusses many variations which may be incorporated into the embodiment illustrated in FIG. 6.” *Id.* at col. 13, ll. 8–10. According to the '273 patent, the “computational tasks” that the LPHDR arithmetic units can perform “enable a variety of practical applications.” *Id.* at col. 17, ll. 18–21. The '273 patent provides, as examples, applications including “finding nearest neighbors,” *id.* at col. 17, l. 29–col. 21, l. 32, “distance weighted scoring,” *id.* at col. 21, l. 34–col. 22, l. 24, and “removing motion blur in images,” *id.* at col. 22, l. 26–col. 23, l. 30.

#### *D. Illustrative Claim*

Challenged claims 1, 33, 36, and 68 of the '273 patent are independent. Claims 2–26, 28, and 32 depend from claim 1, claims 34 and 35 depend from claim 33, claims 37–61, 63, and 67 depend from claim 36, and claims 69 and 70 depend from claim 68. Claim 1 recites:

1. A device:

comprising at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000 and for at least  $X=5\%$  of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least  $X\%$  of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least  $Y=0.05\%$  from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.

IPR2021-00179

Patent 8,407,273 B2

*E. Evidence*

Petitioner relies on the following prior art:

U.S. Patent No. 5,689,677, issued Nov. 18, 1997 (Ex. 1009, “MacMillan”);

U.S. Patent Application Publication No. 2007/0203967 A1, published Aug. 30, 2007 (Ex. 1007, “Dockser”); and

Jonathan Ying Fai Tong, David Nagle, & Rob. A. Rutenbar, “Reducing Power by Optimizing the Necessary Precision/Range of Floating-Point Arithmetic,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 8, no. 3 (June 2000) (Ex. 1008, “Tong”).

*F. Prior Art and Asserted Grounds*

Petitioner asserts that claims 1–26, 28, 32–61, 63, and 67–70 of the ’273 patent are unpatentable on the following grounds:

Claims Challenged	35 U.S.C. §	References/Basis
1, 2, 21–24, 26, 28	103(a) <sup>1</sup>	Dockser
1, 2, 21–24, 26, 28, 32, 33	103(a)	Dockser, Tong
1–26, 28, 36–61, 63	103(a)	Dockser, MacMillan
1–26, 28, 32–61, 63, 67–70	103(a)	Dockser, Tong, MacMillan

## II. ANALYSIS

*A. Level of Ordinary Skill in the Art*

Petitioner asserts that at the time of the earliest possible effective filing date of the ’273 patent (June 19, 2009), a person of ordinary skill in

<sup>1</sup> The Leahy-Smith America Invents Act, Pub. L. No. 112-29, 125 Stat. 284 (2011) (“AIA”), amended 35 U.S.C. § 103. Here, Petitioner’s challenges are based on the pre-AIA version of 35 U.S.C. § 103.

IPR2021-00179

Patent 8,407,273 B2

the art would have had “at least a bachelor’s degree in Electrical Engineering, Computer Engineering, Applied Mathematics, or the equivalent, and at least two years of academic or industry experience in computer architecture.” Pet. 7–8. Patent Owner does not address the level of ordinary skill in the art in its Preliminary Response. Based on the record presented, including our review of the ’273 patent and the types of problems and solutions described in the ’273 patent and cited prior art, we agree with Petitioner’s proposed definition of the level of ordinary skill in the art with one exception. Arguably, the term “at least” creates unnecessary ambiguity. Thus, we delete that term from Petitioner’s definition, and otherwise apply Petitioner’s definition for purposes of this Decision. *See, e.g.*, Ex. 1001, col. 1, l. 26–col. 2, l. 7 (describing in the “Background” section of the ’273 patent various conventional methods of computation and their alleged deficiencies).

### *B. Claim Construction*

We construe the challenged claims

using the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. 282(b), including construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent.

37 C.F.R. § 42.100(b) (2020). We need to construe only those terms “that are in controversy, and only to the extent necessary to resolve the controversy.” *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017). Petitioner states that the terms of the challenged claims should be “given their ordinary and customary meaning as

IPR2021-00179

Patent 8,407,273 B2

understood by a [person of ordinary skill in the art] in accordance with the specification and prosecution history,” but does not propose any express constructions. Pet. 8. Patent Owner disputes the construction of “low precision high dynamic range (LPHDR) execution unit” that Petitioner proposed in the related district court case (i.e., “low precision and high dynamic range processing element designed to perform arithmetic operations on numerical values”), arguing that Dockser does not teach an LPHDR execution unit because “the processing element itself must be fairly characterized as ‘low precision’” and cannot be an execution unit “whose subprecision can be selectively reduced.” Prelim. Resp. 14 (citing Ex. 2001, 13–17) (emphasis omitted). We conclude that no terms require express construction at this time, and address the parties’ arguments regarding whether Dockser teaches the recited LPHDR execution unit below.

*See infra* Section II.C.2.

### *C. Obviousness Ground Based on Dockser*

Petitioner contends that claims 1, 2, 21–24, 26, and 28 are unpatentable over Dockser under 35 U.S.C. § 103(a), citing the testimony of Richard Goodin, P.E., as support. Pet. 8–38 (citing Ex. 1003). We are persuaded that Petitioner has established a reasonable likelihood of prevailing on its asserted ground for the reasons explained below.

#### *1. Dockser*

Dockser discloses performing floating-point operations with a floating-point processor having “selectable subprecision.” Ex. 1007, code (57), ¶¶ 15, 17. “A floating-point representation of a number commonly includes a sign component, an exponent, and a mantissa. To find

IPR2021-00179

Patent 8,407,273 B2

the value of a floating-point number, the mantissa is multiplied by a base (commonly 2 in computers) raised to the power of the exponent. The sign is applied to the resultant value.” *Id.* ¶ 1. “The precision of the floating-point processor is defined by the number of bits used to represent the mantissa. The more bits in the mantissa, the greater the precision.” *Id.* ¶ 2. “The precision of [a] floating-point processor generally depends on the particular application. For example, the ANSI/IEEE-754 standard (commonly followed by modern computers) specifies a 32-bit single format having a 1-bit sign, an 8-bit exponent, and a 23-bit mantissa.” *Id.* “Higher precision results in a higher accuracy, but commonly results in increased power consumption.” *Id.* Dockser explains that the performance of floating-point operations also can be computationally inefficient because “[w]hile some applications may require [high] types of precision, other applications may not.” *Id.* ¶ 3. For example, “some graphics applications may only require a 16-bit mantissa,” such that “any accuracy beyond 16 bits of precision tends to result in unnecessary power consumption,” but other applications may require “greater precision.” *Id.* Accordingly, there was “a need in the art for a floating-point processor in which the reduced precision, or subprecision, of the floating-point format is selectable.” *Id.*

Figure 1 of Dockser is reproduced below.

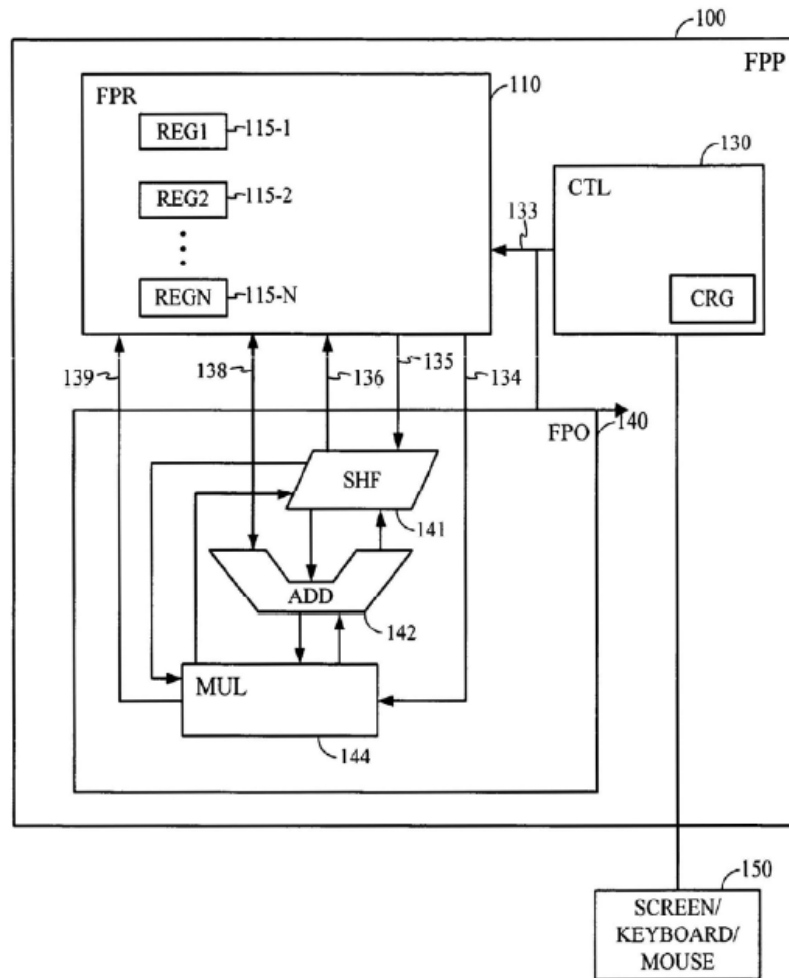


FIG. 1

Figure 1 depicts floating-point processor (FPP) 100 including floating-point register file (FPR) 110 for storing floating-point numbers, floating-point controller (CTL) 130 “used to select the subprecision of the floating-point operations using a control signal 133,” and floating-point mathematical operator (FPO) 140 with components “configured to perform the floating-point operations,” such as floating-point adder (ADD) 142 and floating-point multiplier (MUL) 144. *Id.* ¶¶ 15, 18, 19.

Figure 2 of Dockser is reproduced below.

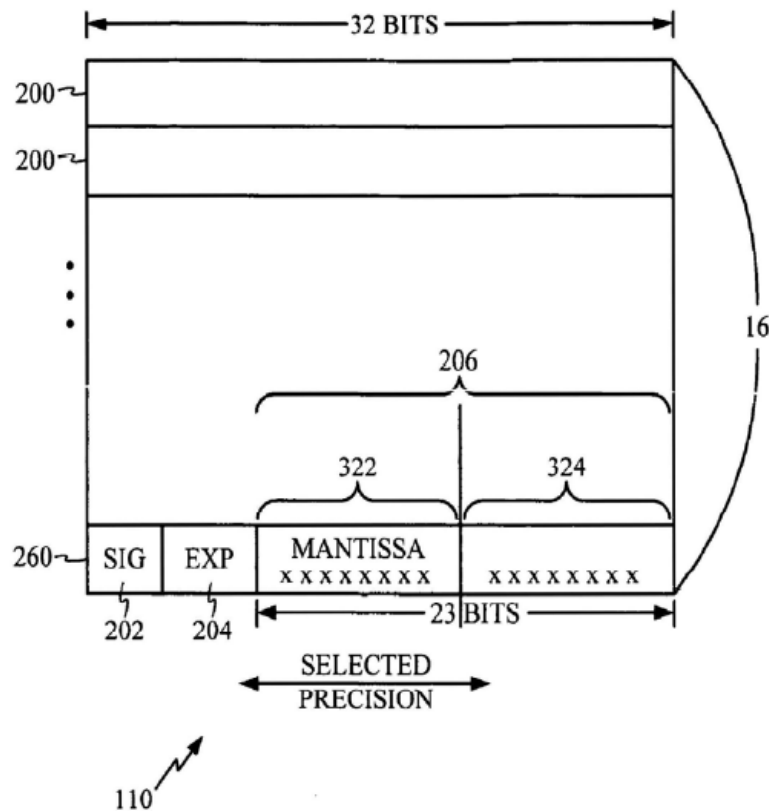


FIG. 2

Figure 2 depicts an exemplary data structure for floating-point register file 110 including 16 addressable register locations 200, each “configured to store a 32-bit binary floating-point number” as “a 1-bit sign 202, an 8-bit exponent 204, and a [23-bit] fraction 206.” *Id.* ¶ 17.

“[F]or each instruction of a requested floating-point operation, the relevant computational unit . . . receive[s] from the floating-point register file 110 one or more operands stored in one or more of the register locations” and executes the instruction “at the subprecision selected by the floating-point controller 130.” *Id.* ¶¶ 23–24. The precision of the floating-point operation can be reduced by “caus[ing] power to be removed from the floating-point register elements for the excess bits of the fraction that are not required to meet the precision specified by the subprecision

IPR2021-00179

Patent 8,407,273 B2

select bits” written to the control register. *Id.* ¶¶ 6, 25–26. For example, “if each location in the floating-point register file contains a 23-bit fraction, and the subprecision required for the floating-point operation is 10-bits, only the 9 commonly significant bits (MSBs) of the fraction are required; the hidden or integer bit makes the tenth.” *Id.* ¶ 26. “Power can be removed from the floating-point register elements for the remaining 14 fraction bits.” *Id.*

Alternatively, power can be removed in elements of “the logic in the floating-point operator 140 that remains unused as a result of the subprecision selected.” *Id.* ¶¶ 7, 27, 29, 32, Fig. 2 (depicting mantissa fraction 206 as having portion 322 for powered bits and portion 324 for unpowered bits). Figures 3A and 3B of Dockser show such removal of power to the floating-point operator logic for a floating-point addition and floating-point multiplication operation, respectively. Figure 3B of Dockser is reproduced below.

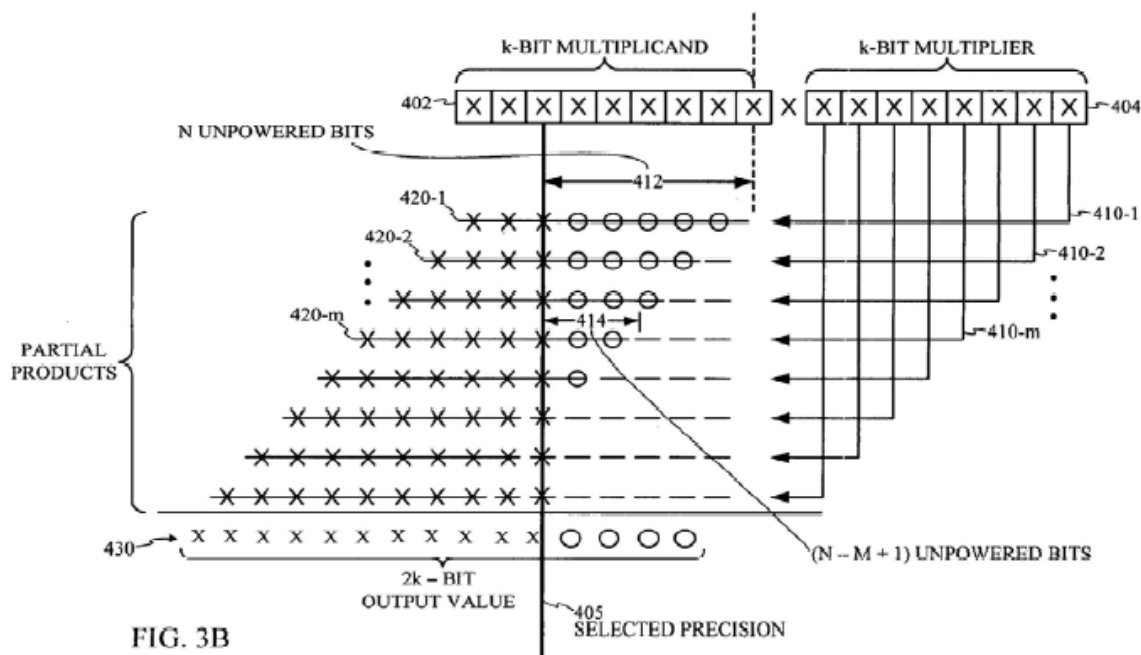


Figure 3B depicts k-bit multiplicand 402 and k-bit multiplier 404 to be multiplied together “using a shift-and-add technique,” where the

multiplication occurs in stages 410-1 through 410-m. *Id.* ¶¶ 30–31. A partial product 420-i is generated for every bit in multiplier 404 at corresponding stage 410-i and then left-shifted “as a function of the multiplier bit with which it is associated, after which the operation moves on to the next stage.” *Id.* ¶ 31. The partial products are eventually added together to generate output value 430. *Id.* “[P]ower may be removed from the logic used to implement the stages to the right of the line 405” indicating the selected subprecision. *Id.* ¶¶ 32–33.

## 2. Claim 1

### a) *Petitioner’s Mapping of Dockser to the Limitations of Claim 1*

Petitioner argues that Dockser teaches or suggests all of the limitations of claim 1. Pet. 8–36. Petitioner contends that Dockser teaches a “device” (i.e., computing system) comprising a “low precision high dynamic range (LPHDR) execution unit” (i.e., the FPP).<sup>2</sup> *Id.* at 12–13. With respect to the “low precision” aspect of the limitation, Petitioner argues that the FPP is “low precision” because “‘the precision’ of operations in the FPP is ‘reduced’” and because the FPP “operates with the minimum imprecision” required by the subsequent language in claim 1. *Id.* at 13–14 (quoting Ex. 1007 ¶ 14). With respect to the “high range” aspect of the limitation, Petitioner asserts that the FPP “uses an 8-bit floating-point exponent . . . that provides an even higher dynamic range” than the 6-bit floating-point

---

<sup>2</sup> Petitioner also provides an “alternative mapping” where “the floating-point operator (FPO) inside Dockser’s FPP” constitutes an LPHDR execution unit. Pet. 13, 18–20, 35. We need not evaluate those arguments at this time, as we determine that Petitioner has made a sufficient showing on the current record that the FPP is an LPHDR execution unit.

IPR2021-00179

Patent 8,407,273 B2

exponent disclosed in the '273 patent. *Id.* at 14 (citing Ex. 1007 ¶ 17; Ex. 1001, col. 14, ll. 53–61).

Petitioner argues that Dockser's FPP is adapted to execute a "first operation" (e.g., "reduced-precision multiplication") on a "first input signal representing a first numerical value" (i.e., input electrical signal representing an operand received at the registers) to produce a "first output signal representing a second numerical value" (i.e., output electrical signal representing an operand sent to a register and then main memory), where the FPP "performs operations on the[] inputs via the FPP's data paths 134–139 and components 140–144." *Id.* at 14–18. Petitioner further argues that "the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/65,000 through 65,000," as recited in claim 1, because the FPP "operates on IEEE-754 32-bit single-format numbers having 8-bit exponents" and, therefore, the dynamic range of normal operands would be "from around  $2^{-126}$  (much smaller than 1/65,000) to around  $2^{127}$  (much larger than 65,000)." *Id.* at 19.

The final limitation of claim 1 is that

for at least  $X=5\%$  of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least  $X\%$  of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least  $Y=0.05\%$  from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.

We refer to this as "the imprecision limitation." With respect to the "statistical mean" aspect of the imprecision limitation, Petitioner argues that a person of ordinary skill in the art would have understood the limitation

IPR2021-00179

Patent 8,407,273 B2

in the context of the '273 patent's stated intent to claim not only "repeatable" deterministic embodiments like digital circuits that always produce the same output when repeating an operation on the same input, but also analog embodiments that are non-deterministic because they "introduce noise into their computations, so the computations are not repeatable."

*Id.* at 20–21 (quoting Ex. 1001, col. 4, ll. 7–13). Specifically, for "non-deterministic embodiments," the statistical mean would be the average of "the different outputs produced by the same operation on the same input." *Id.* at 21. For "deterministic digital embodiments" like Dockser, though, the statistical mean is "the same as the numerical value of the first output signal for any individual execution of the first operation on each specific input, because that output is always the same for any specific input." *Id.* Consequently, repeatedly executing a multiplication operation using Dockser's floating-point multiplier "on the same input (*i.e.*, pair of operands) with the same precision level yields the same result for every execution; therefore, the statistical mean of the outputs is the same as the output for any single execution." *Id.* at 22.

With respect to the "exact mathematical calculation" aspect of the imprecision limitation, Petitioner argues that because Dockser performs a reduced-precision multiplication, the result of the operation differs from what would be the exact mathematical result of the operation, namely "the (>32-bit) product that would result if the pair of input 32-bit operands were multiplied without reducing precision." *Id.* at 22–23 (emphasis omitted).

Petitioner further explains how Dockser teaches that "for at least X=5% of the possible valid inputs to the first operation," the statistical mean of the results of executing the first operation "differs by at least Y=0.05%" from the result of the exact mathematical calculation. *Id.* at 23–36. The

“possible valid inputs” in Dockser are “the set of possible normal IEEE-754 32-bit single-format numbers forming pairs of operands in input signals to the execution unit that can be multiplied together to produce an output representing a numerical value (rather than, *e.g.*, an overflow/underflow exception).” *Id.* at 23. Petitioner argues that Dockser’s FPP operates at a precision level meeting the claimed X and Y percentages for such input pairs, pointing to Dockser’s description of retaining only some of the bits of a mantissa fraction (*e.g.*, the 9 most-significant bits of a 23-bit fraction, as shown in Figure 2 above) and dropping the remaining “excess bits” (*e.g.*, the other 14 bits). *Id.* at 23–24. Dockser teaches dropping the excess bits to reduce precision by either (1) “removing power from storage elements in the FPP’s registers that correspond to the excess (dropped) mantissa bits,” or (2) “removing power from elements within the multiplier logic that computes the product of the operand mantissas.” *Id.* at 25. Petitioner contends that both techniques, used either individually or in combination, teach the recited imprecision. *Id.* at 26–36. We determine that Petitioner has made a sufficient showing on the current record with respect to the first precision-reducing technique, and need not evaluate Petitioner’s alternative arguments at this time.

Relying on the first technique and Dockser’s example of retaining 9 mantissa bits and dropping 14 mantissa bits, Petitioner argues that a person of ordinary skill in the art “would have understood the output of unpowered storage elements would be tied to zero voltage (*e.g.*, ground), making those 14 ‘excess’ bits zeroes.” *Id.* at 28. Multiplying two operands with excess bits dropped results in an output with reduced precision from the exact product, where the amount of error depends on the number of mantissa bits dropped. *Id.* at 29–30. Petitioner contends that an ordinarily skilled artisan

IPR2021-00179

Patent 8,407,273 B2

would have understood, by “straightforward math” described in Appendix I.A to the Petition, that “the relative error (the claimed ‘Y’ percentage) of any floating-point number output from Dockser’s reduced-precision multiplication is the same as the relative error of its mantissa, independent of its exponent and sign.” *Id.* at 29.

Petitioner provides a detailed explanation as to why a person of ordinary skill in the art would have understood Dockser as teaching the recited X and Y percentages of claim 1. *Id.* at 30–32. First, Petitioner states that “[g]iven the massive number of possible inputs to Dockser’s FPP (including over 70 trillion possible pairs of normal IEEE-754 single-format mantissas), a [person of ordinary skill in the art] would have performed Dockser’s FPP operation in software to determine the fraction X of all possible valid inputs that produce at least the claimed relative error Y when a given number of mantissa bits are dropped.” *Id.* at 30–31. Petitioner states that Mr. Goodin wrote such a program to perform reduced-precision multiplication retaining 9 mantissa bits and dropping the 14 excess bits as in Dockser that tested all possible valid mantissa pairs and “produce[d] at least Y=0.05% relative error for 92.1% of possible valid inputs (greater than X=5%).” *Id.* at 31, 68–70 (citing Ex. 1003 ¶¶ 290–292, 441–455).

Second, Petitioner argues that a person of ordinary skill in the art “would also have understood algebraically that Dockser’s register bit-dropping technique meets [the limitation] by examining the absolute *minimum* relative error produced by zeroing certain mantissa bit positions.” *Id.* at 31. Petitioner explains how over 12% of possible input operands “have a zero as their most-significant (leftmost) mantissa fraction bit and ones as their tenth and eleventh fraction bits,” such that retaining only 9 mantissa bits of operands would result in “every input in that

12% produc[ing] at *minimum* 0.097% relative error.” *Id.* at 31–32, 71–76 (citing Ex. 1003 ¶¶ 293–294, 456–475). Petitioner’s contentions regarding claim 1 are supported by the testimony of Mr. Goodin. *See id.* at 8–36; Ex. 1003 ¶¶ 193–294, 441–475.

*b) Patent Owner’s Arguments*

Patent Owner makes two arguments in its Preliminary Response disputing Petitioner’s contentions regarding claim 1. Prelim. Resp. 11–24.

*(1) LPHDR Execution Unit*

First, Patent Owner argues that Dockser’s FPP is not a “low precision high dynamic range (LPHDR) execution unit” because it is also “capable of operating at full precision.” Prelim. Resp. 11–12. According to Patent Owner, “Dockser discloses a 32-bit FPP that includes all of the circuitry needed for full precision arithmetic on data in IEEE 32-bit format, and also having additional circuitry allowing for selectable subprecisions.” *Id.* at 11. As support for this reading of Dockser, Patent Owner points to statements in Dockser that precision “may” be reduced and that certain applications require “greater precision.” *Id.* at 11–12 (citing Ex. 1007 ¶¶ 3, 14, 26, 28). Patent Owner also argues that Petitioner incorrectly focuses on whether the FPP is “capable of performing a few operations that are ‘low precision’” in a “9-bit subprecision mode,” rather than “whether the execution unit itself can be fairly characterized as being ‘low precision.’” *Id.* at 13 (emphasis omitted). Patent Owner asserts that Dockser’s FPP is “a full precision processor that—even when performing operations in a reduced precision mode—includes all of the circuitry and capability needed to perform full precision operations.” *Id.*

We disagree based on the current record. Patent Owner's position is that, to be an LPHDR execution unit, the execution unit must be capable of low-precision operations and nothing else. We are not persuaded that the claim language is so limiting. Claim 1 recites a "device . . . comprising at least one first low precision high dynamic range (LPHDR) execution unit." The only limitations on the execution unit recited in the claim are that the execution unit be "low precision," "high dynamic range," and "adapted to execute a first operation" meeting certain criteria specified in the imprecision limitation (i.e., a minimum relative error Y for a minimum fraction X of possible valid inputs in a specified dynamic range). As Petitioner points out, the claim does not recite any structural characteristics of the execution unit and does not include any negative limitation precluding the execution unit from performing other types of operations. *See* Pet. 30; Reply 1. The recitation of "a first operation" in claim 1 further requires only one or more first operations (that meet the low precision criteria specified in the claim); it does not say that "every" operation must be low precision or exclude other capabilities for other operations. *See KCJ Corp. v. Kinetic Concepts, Inc.*, 223 F.3d 1351, 1356 (Fed. Cir. 2000) ("This court has repeatedly emphasized that an indefinite article 'a' or 'an' in patent parlance carries the meaning of 'one or more' in open-ended claims containing the transitional phrase 'comprising.'"). On this record, we see no reason why the claim precludes the execution unit from having additional circuitry for performing other types of operations, as long as the execution unit is capable of performing the recited first operation and meeting the criteria set forth in the imprecision limitation.

Petitioner has provided sufficient evidence at this stage that Dockser performs operations at the precision level specified in the imprecision

IPR2021-00179

Patent 8,407,273 B2

limitation. Dockser’s FPP “perform[s] certain mathematical operations,” such as “multiplication,” at “reduced precision.” Ex. 1007 ¶¶ 1, 30–32. Petitioner expressly identifies “reduced-precision multiplication” as the alleged “first operation” in Dockser. Pet. 14, 16–17; *see also* Prelim. Resp. 10 (acknowledging that Dockser’s FPP performs “a reduced-precision operation”), 11 (acknowledging that the FPP has “additional circuitry [that] allows the FPP to perform some operations with selectively reduced precision”). Petitioner further explains how it calculated the relative error for multiplication of all possible input operands when the operands have 9 retained mantissa bits and 14 dropped mantissa bits.<sup>3</sup> Pet. 24, 28–32. That is consistent with Dockser, which discloses a specific example of multiplying 23-bit operands having 9 retained mantissa bits and 14 dropped mantissa bits. *See* Ex. 1007 ¶¶ 26, 29.

Importantly, as Patent Owner acknowledges, Dockser’s example of 9-bit precision multiplication is a distinct operation from multiplication at other levels of precision. *See* Prelim. Resp. 21 (arguing that Dockser can “perform an addition operation  $a+b$  in full precision mode, 22-bit mode, 21-bit mode, etc.” and that “[e]ach of these operations is distinct, and generally produces different outputs, with some being more precise than others”). Dockser describes exactly how such multiplication is performed, and the functionality of the FPP differs depending on which level of precision multiplication is selected. *See, e.g.*, Ex. 1007 ¶¶ 30–34, Fig. 3B;

---

<sup>3</sup> Accordingly, we do not agree with Patent Owner that Petitioner improperly made a new argument in its Reply that “multiplying two input values in a way that reduces precision down to 9 mantissa bits” constitutes a “first operation.” Sur-Reply 2.

Pet. 23–32. On this record, we find that the fact that Dockser is capable of performing other, different operations (e.g., 23-bit full precision) does not detract from Dockser’s disclosure of a specific example meeting the “low precision” requirements of the claim.

Patent Owner further argues that Petitioner fails to show that it would have been obvious to modify Dockser’s FPP to be an LPHDR execution unit. Prelim. Resp. 14–17. Patent Owner contends that Dockser teaches away from such a unit because

rather than committing itself to performing low precision operations on a very high percentage of all possible valid inputs (in order to shrink the size of the execution unit), Dockser specifically teaches that its selectable precision (which means supporting full precision and programmability, both of which increase the size of the execution unit) is the key feature of the Dockser FPP.

*Id.* at 15 (citing Ex. 1007 ¶ 3). Dockser’s FPP “is more complex than conventional 32-bit microprocessors, not less,” so an ordinarily skilled artisan would not be motivated to “reverse course” and make it only capable of low precision operations according to Patent Owner. *Id.* at 17 (emphasis omitted). As explained above, however, we disagree on this record with Patent Owner’s position that an LPHDR execution unit must be incapable of anything other than low precision operations, and thus do not agree that a modification to Dockser in that respect would have been necessary. We find Petitioner’s analysis as to why a person of ordinary skill in the art would have understood Dockser’s FPP to be an LPHDR execution unit, supported by the testimony of Mr. Goodin, sufficient at this early stage.

(2) *Imprecision Limitation*

Second, Patent Owner argues that Dockser does not teach or suggest the imprecision limitation because Petitioner fails to show that Dockser’s output over all “possible valid inputs” meets the limitation. Prelim. Resp. 18–24 (emphasis omitted). According to Patent Owner, “[t]he range of possible valid inputs for Dockser” includes all possible IEEE 32-bit values “across the entire range of the possible subprecision select bits,” not just operands with 9 retained mantissa bits and 14 dropped mantissa bits as Petitioner discussed in the Petition. *Id.* at 19–20. Patent Owner argues that Mr. Goodin’s software program and algebraic explanation address “only a small subset of Dockser’s possible valid inputs” because they fail to account for “whether the imprecision limitation is met for subprecision select bits corresponding to any implementation where fewer than 14 bits are dropped.” *Id.* at 20–22. In Patent Owner’s view, the analysis of Dockser needs to take into account operations at “full precision mode, 22-bit mode, 21-bit mode, etc.” *Id.* at 21.

We disagree based on the current record. The imprecision limitation recites that “for at least  $X=5\%$  of the possible valid inputs to *the first operation*, the statistical mean, over repeated execution of *the first operation* on each specific input from the at least  $X\%$  of the possible valid inputs to *the first operation*,” “executing *the first operation* on that input differs by at least  $Y=0.05\%$  from the result of an exact mathematical calculation of *the first operation* on the numerical values of that same input” (emphasis added). The relevant inquiry thus looks at all possible valid inputs to the first operation—not all possible valid inputs to the LPHDR execution unit overall, or all possible valid inputs to other operations. Again, in its mapping of Dockser to claim 1, Petitioner expressly identifies

“reduced-precision multiplication” as “the first operation.” Pet. 16–17. Dockser discloses a specific example of multiplying operands having 9 retained mantissa bits and 14 dropped mantissa bits. Ex. 1007 ¶¶ 26, 30–34, Fig. 3B. The relevant possible valid inputs for assessing the imprecision limitation, therefore, are the inputs to that operation. On this record, we see no reason why the analysis of Dockser also needs to account for other operations, such as “full precision” multiplication, which Patent Owner acknowledges are “distinct” operations. *See* Prelim. Resp. 21.

Patent Owner also challenges Mr. Goodin’s analysis as “impermissibly applying hindsight and using the claims as a roadmap” because “he opines that a [person of ordinary skill in the art] would operate Dockser with the *express goal* of dropping enough bits from the mantissa to meet the imprecision limitation.” *Id.* at 16 (citing Ex. 1003 ¶¶ 288–289). We disagree based on the current record. The precision level of 9 retained mantissa bits and 14 dropped mantissa bits that Mr. Goodin analyzes is not chosen with the goal of meeting the imprecision limitation, but rather is the specific example described in Dockser. *See* Ex. 1003 ¶¶ 263, 277–278, 282, 284. Mr. Goodin describes the calculations performed by his software program and provides an algebraic explanation for why he believes that a person of ordinary skill in the art would have understood Dockser’s description of multiplication at the 9-bit precision level teaches the imprecision limitation. *See id.* ¶¶ 290–294, 441–475. Further, Mr. Goodin cites extensively to language in Dockser describing reduced-precision multiplication in support of his opinions. *See id.* ¶¶ 256–267. Patent Owner has not presented any evidence at this stage indicating that Mr. Goodin’s calculations for Dockser’s 9-bit precision level are factually incorrect. We have reviewed that supporting testimony and are persuaded, based on the

current record, that a person of ordinary skill in the art would have understood Dockser to teach the imprecision limitation of claim 1.

*c) Conclusion*

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claim 1 is unpatentable over Dockser.

*3. Claims 2, 21–24, 26, and 28*

We have reviewed Petitioner’s contentions regarding claims 2, 21–24, 26, and 28, which depend from claim 1, and are persuaded that Petitioner has made a sufficient showing at this stage for those claims as well. *See* Pet. 36–38. Petitioner explains how each limitation of the dependent claims is taught or rendered obvious by the disclosure of Dockser. *Id.* For example, claim 2 recites that “the at least one first LPHDR execution unit comprises at least part of an FPGA.” Dockser discloses a FPP that is part of “a field programmable gate array (FPGA) or other programmable logic component.” Ex. 1007 ¶ 35; *see* Pet. 36. Petitioner’s contentions are supported by the testimony of Mr. Goodin and are persuasive based on the current record. *See* Pet. 36–38; Ex. 1003 ¶¶ 311–322. Patent Owner does not argue the challenged dependent claims separately, only disputing Petitioner’s arguments regarding independent claim 1. *See* Prelim. Resp. 8–24. We disagree with Patent Owner’s arguments based on the current record for the reasons explained above. *See supra* Section II.C.2.b.

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 2, 21–24, 26, and 28 are unpatentable over Dockser.

*D. Obviousness Ground Based on Dockser and Tong*

Petitioner contends that claims 1, 2, 21–24, 26, 28, 32, and 33 are unpatentable over Dockser and Tong under 35 U.S.C. § 103(a), citing the testimony of Mr. Goodin as support. Pet. 38–45 (citing Ex. 1003). We are persuaded that Petitioner has established a reasonable likelihood of prevailing on its asserted ground for the reasons explained below.

*1. Tong*

Tong is an IEEE journal article entitled “Reducing Power by Optimizing the Necessary Precision/Range of Floating-Point Arithmetic.” Ex. 1008, 273. Tong teaches reducing power consumption by minimizing the bitwidth representation of floating-point data. *Id.* According to Tong, using a variable bitwidth floating-point unit saves power. *Id.*

*2. Claim 1*

Petitioner asserts that “Tong, like Dockser,” “confirms that the number of mantissa bits used in a high-dynamic-range floating-point execution unit was a well-known result-effective variable impacting power consumption and precision.”<sup>4</sup> Pet. 39–40 (citing Ex. 1008, 273–278; Ex. 1003 ¶¶ 324–325).

---

<sup>4</sup> Petitioner presents sufficient evidence on the current record to establish a reasonable likelihood that Tong is a prior art printed publication under 35 U.S.C. § 102(b). *See* Pet. 38–39 (citing Ex. 1025 ¶¶ 8–11; Ex. 1026, 27; Ex. 1027, 27).

IPR2021-00179

Patent 8,407,273 B2

Petitioner relies on Tong's Figure 6, reproduced below with Petitioner's annotations. *Id.* at 40.

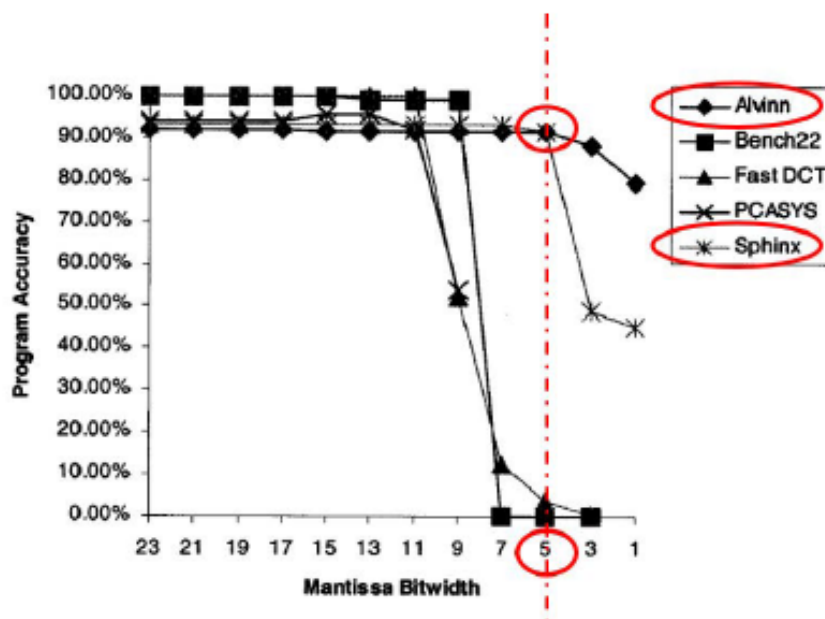


Fig. 6. Program accuracy across various mantissa bitwidths.

Tong's Figure 6, above, is a line graph showing program accuracy, from 0% to 100%, on the vertical axis and mantissa bitwidth, from 1 to 23, on the horizontal axis. Ex. 1008, 279. The figure shows this data for five programs: ALVINN, Bench22, Fast DCT, PCASYS, and Sphinx. *Id.* The programs implement different signal-processing tasks. *Id.* at 278 (Table IV). ALVINN, for example, is a neural network trainer that uses backpropagation. *Id.* And Sphinx is a speech-recognition program. *Id.*

For the ALVINN and Sphinx line plots in Tong's Figure 6, Petitioner adds a dashed red line extending vertically through the data points with a mantissa bitwidth of 5. Pet. 40. Petitioner asserts that, for these programs, Tong teaches that "the accuracy does not change significantly with as few as 5 mantissa [fraction] bits." *Id.* (quoting Ex. 1008, 278, Section V.B). In Petitioner's view, Tong omits unnecessary bits to reduce waste and power consumption. *Id.* at 41. Petitioner characterizes Tong as "[h]aving

empirically determined the minimum number of mantissa bits necessary to maintain acceptable accuracy of particular applications.” *Id.* (citing Ex. 1008, 273, 274, 279, 284).

Petitioner concludes that “Tong’s teaching that retaining 5 mantissa fraction bits is sufficient in some applications (including ALVINN and Sphinx[]) would have motivated a [person of ordinary skill in the art] to configure Dockser’s FPP . . . to operate at a selected precision level retaining as few as 5 mantissa fraction bits.” *Id.* at 41–42. According to Petitioner, one of ordinary skill in the art would have done so “to conserve power when running those applications, or others empirically determined to not require greater precision using Tong’s techniques.” *Id.* at 42 (citing Ex. 1003 ¶ 334). Petitioner also concludes that “[d]etermining the optimum range of imprecision to achieve the best power reduction without sacrificing accuracy for a particular application was a matter of routine optimization of a result-effective variable.” *Id.* at 43 (citing Ex. 1003 ¶ 339).

In response, Patent Owner argues that Dockser is deficient for the same reasons discussed in connection with the challenge based on Dockser alone and Tong does not remedy those deficiencies. Prelim. Resp. 24–25. In Patent Owner’s view, Petitioner’s Dockser-Tong analysis, like the analysis of Dockser alone, accounts for “only a cherry-picked subset of the ‘possible valid inputs.’” *Id.* at 24 (citing Pet. 44–45). Patent Owner argues that Petitioner’s analysis ignores the remaining possible valid inputs across all subprecision selections. *Id.* at 24–25.

Patent Owner’s argument about all subprecision selections, however, does not squarely address Petitioner’s obviousness rationale—i.e., that one of ordinary skill in the art would have been motivated to configure Dockser’s FPP at a “selected precision level” according to Tong. *See*

Pet. 41–42. Specifically, Petitioner’s rationale proposes using a selected precision level “retaining as few as 5 mantissa fraction bits.” *Id.* Instead of addressing the analysis of Dockser at this selected precision level, Patent Owner’s arguments focus on Dockser alone. *See* Prelim. Resp. 24–25.

Petitioner’s relative-error analysis of the Dockser-Tong combination, though, is different from the one in the ground based on Dockser alone. *See* Pet. 42. In particular, Petitioner used a software program to determine the relative error when retaining 5 mantissa fraction bits—i.e., the number of bits that Dockser would use under Petitioner’s proposed combination with Tong. *Id.* (citing Ex. 1003 ¶ 336). Also, the algebraic analysis cited by Petitioner uses 5 mantissa fraction bits. *Id.* (citing Ex. 1003 ¶ 336). Because Petitioner’s rationale relies on modifying Dockser based on Tong to have a particular precision level, we are not persuaded at this stage that Petitioner’s relative-error analysis is deficient for not considering all possible subprecision selections. *See* Prelim. Resp. 24–25; *supra* Section II.C.2.b.(2) (explaining why Petitioner’s asserted ground based on Dockser alone is not deficient for failing to address all possible IEEE 32-bit values, rather than just operands with 9 retained mantissa bits and 14 dropped mantissa bits).

We find Petitioner’s analysis, supported by the testimony of Mr. Goodin, sufficient to show a reasonable likelihood of prevailing on its assertion that claim 1 is unpatentable over Dockser and Tong.

### 3. *Claims 2, 21–24, 26, 28, 32, and 33*

We have reviewed Petitioner’s contentions regarding claims 2, 21–24, 26, 28, and 32, which depend from claim 1, as well as independent claim 33, which recites an LPHDR execution unit and the imprecision limitation, and are persuaded that Petitioner has made a sufficient showing at this stage for

those claims as well. *See* Pet. 36–38, 41–45. Petitioner explains how each limitation of the claims is taught or rendered obvious by the disclosure of Dockser or the combination of Dockser and Tong. *Id.* Petitioner’s contentions are supported by the testimony of Mr. Goodin and are persuasive based on the current record. *See id.*; Ex. 1003 ¶¶ 311–322, 333–348.

Patent Owner does not argue the challenged dependent claims separately, only disputing Petitioner’s arguments regarding the challenged independent claims. *See* Prelim. Resp. 8–25. Based on the current record, we disagree with Patent Owner’s arguments regarding the LPHDR execution unit and imprecision limitation for the reasons explained above. *See supra* Section II.C.2.b.

Patent Owner makes an additional argument regarding independent claim 33, which recites a “device comprising a computer processor and a computer-readable memory storing computer program instructions, wherein the computer program instructions are executable by the processor to emulate a second device comprising” an LPHDR execution unit having the same properties as recited in claim 1. Patent Owner contends that “a person of ordinary skill would not have been motivated to incorporate Tong’s ‘emulation’ into the systems of Dockser” because “Tong uses software emulation only for the purposes of *investigating* the behavior of *physical* devices.” Prelim. Resp. 25 (citing Ex. 1008, 273, 278). We disagree based on the current record. Petitioner’s position is that “Tong’s teaching to ‘emulate[] in software different bitwidth FP units’ ‘to determine application accuracy’ . . . would have motivated a [person of ordinary skill in the art] to *emulate the Dockser/Tong device . . . in software* to assess the accuracy of applications running on the device at selected precision levels.” Pet. 44

(quoting Ex. 1008, 278; citing Ex. 1003 ¶¶ 341–342) (emphasis added). Emulating the device of Dockser (modified based on Tong) in software appears, on this record, to be consistent with Tong, which describes the emulation of a floating-point (FP) device to determine “the relationship between program accuracy and number of bits in FP representation.” See Ex. 1008, 278. Based on the current record, we do not see why Tong’s disclosures regarding FP hardware would have dissuaded an ordinarily skilled artisan from making the combination that Petitioner proposes.

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 2, 21–24, 26, 28, 32, and 33 are unpatentable over Dockser and Tong.

*E. Obviousness Ground Based on Dockser and MacMillan*

Petitioner contends that claims 1–26, 28, 36–61, and 63 are unpatentable over Dockser and MacMillan under 35 U.S.C. § 103(a), citing the testimony of Mr. Goodin as support. Pet. 45–56 (citing Ex. 1003). We are persuaded that Petitioner has established a reasonable likelihood of prevailing on its asserted ground for the reasons explained below.

*1. MacMillan*

MacMillan is entitled “Circuit for Enhancing Performance of a Computer for Personal Use.” Ex. 1009, code (54). MacMillan teaches using Single Instruction Multiple Data (SIMD) parallel-processing architectures for adding supercomputer performance to personal-use computers. *Id.* at col. 5, ll. 22–54. MacMillan’s computer system comprises a “Host CPU” (i.e., “a 386, 486 or Pentium[] processor”) and SIMD-random access memory (SIMD-RAM) device. *Id.* at col. 9, ll. 30–31, Figs. 3, 5.

MacMillan describes an example architecture where the SIMD-RAM device has 256 processing elements (PEs), but states that the disclosed architecture “allows scaling to higher or lower density chips with more or fewer PEs.” *Id.* at col. 12, l. 60–col. 13, l. 4, col. 13, ll. 38–41, col. 16, ll. 20–22.

## 2. Claim 1

In this asserted ground, Petitioner relies on MacMillan for its teachings about multiple floating-point execution units and concludes that it would have been obvious to implement a device with multiple Dockser FPPs operating in parallel. Pet. 47–49.

Patent Owner argues that MacMillan does not remedy Dockser’s deficiencies. Prelim. Resp. 25–28. Apart from this argument, Patent Owner does not present arguments specifically analyzing MacMillan as it is used in Petitioner’s challenge to claim 1. *Id.* Rather, Patent Owner refers to the arguments regarding the challenge based on Dockser alone. *Id.* at 25–26. Thus, for the reasons discussed in connection with the asserted ground based on Dockser alone, and because we conclude that Petitioner has made a sufficient showing of obviousness based on the combination with MacMillan, we determine on this record that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claim 1 is unpatentable over Dockser and MacMillan. *See supra* Section II.C.2.b.

## 3. Claims 5, 8, 10, 40, 43, and 45

Claim 5 recites that “the number of LPHDR execution units in the device exceeds by at least ten the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.” Claims 8 and 10

IPR2021-00179

Patent 8,407,273 B2

recite the same limitation, replacing “at least ten” with “at least one hundred” and “at least five hundred,” respectively. Claims 40, 43, and 45, which depend from claim 36, mirror claims 5, 8, and 10.

With respect to these claims, Petitioner asserts that MacMillan teaches 256 processing elements (PEs), and the Dockser-MacMillan combination would have a “single Host CPU floating-point unit” and at least one FPP in each of the PEs, of which there can be “256 or more, up to ‘tens of thousands.’” Pet. 49–51, 56 (quoting Ex. 1009, col. 2, ll. 13–15; citing Ex. 1003 ¶¶ 361, 369, 370, 374). Thus, in Petitioner’s proposed combined device based on Dockser and MacMillan, the number of LPHDR execution units (Dockser FPPs) “exceeds its number (one) of traditional-precision execution units (the single Host CPU floating-point unit) by more than 10, more than 100, and more than 500.” *Id.* at 50–51 (citing Ex. 1003 ¶¶ 369–370, 374).

Patent Owner argues that “Dockser and MacMillan cannot possibly disclose the limitations . . . even if one were to assume the Dockser FPP to be an LPHDR execution unit.” Prelim. Resp. 27. Patent Owner argues that, because (1) Petitioner asserts that MacMillan’s processor is the only execution unit (EU) for purposes of the claims, and (2) Dockser’s FPPs are designed to perform multiplication on 32-bit numbers, “the number of traditional precision units in the Dockser/MacMillan combination must *always* be greater than or equal to the number of ‘LPHDR’ Dockser FPPs.” *Id.*

But, under Petitioner’s theory, the identified 32-bit EUs in claims 5, 8, 10, 40, 43, and 45 (i.e., “execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide”) are different from what Petitioner identifies as the

IPR2021-00179

Patent 8,407,273 B2

LPHDR execution units (i.e., a Dockser FPP in each of the 256 (or more) PEs of MacMillan). *See* Pet. 49–51. So, under Petitioner’s rationale, Dockser’s FPPs do not count toward the number of 32-bit EUs in the Dockser-MacMillan combination. *See id.* Specifically, Petitioner has shown, on this preliminary record, that the claimed “low precision high-dynamic range (LPHDR) execution unit” at least encompasses Dockser’s FPP. *See supra* Section II.C.2.b.(1). Petitioner asserts that the ’273 patent describes the LPHDR execution units as “sometimes” producing results that are different from the correct result, in contrast to 32-bit EUs, which the ’273 patent describes as “‘traditional precision’ execution units that do not ‘sometimes’ produce results different from the correct traditional-precision result.” Pet. 50. Thus, if the proposed combination has 256 (or more) FPPs (LPHDR execution units) and one host CPU (the 32-bit EUs), the number of LPHDR execution units would be greater than the number of 32-bit EUs by the amounts recited in claims 5, 8, 10, 40, 43, and 45.

Petitioner supports its contentions with respect to claims 5, 8, 10, 40, 43, and 45 with citations to the ’273 patent. We preliminarily agree with Petitioner that the cited parts of the patent distinguish between the two sets of units. For example, the patent describes the 32-bit arithmetic elements as “traditional” precision: “‘arithmetic elements . . . designed to perform . . . floating point arithmetic with a word length of 32 or more bits’ are ‘designed to perform . . . arithmetic of traditional precision.’” *Id.* (citing Ex. 1001, col. 27, l. 63–col. 28, l. 3). As for the LPHDR units, the paragraph cited by Petitioner explains how the precision may vary across implementations. *See* Ex. 1001, col. 26, l. 50–col. 27, l. 4. We preliminarily agree with Petitioner that all of these embodiments, unlike the 32-bit arithmetic elements, are

described as producing results that are “sometimes” or “all of the time” not closer than a certain amount to the correct result. *See* Pet. 50 (citing Ex. 1001, col. 26, ll. 50–60). Thus, Petitioner’s distinction between the 256 (or more) FPPs and one host CPU in the proposed combination is consistent with the language of the challenged claims and adequately supported on this preliminary record.

We find Petitioner’s analysis, supported by the testimony of Mr. Goodin, sufficient to show a reasonable likelihood of prevailing on its assertion that claims 5, 8, 10, 40, 43, and 45 are unpatentable over Dockser and MacMillan.

*4. Claims 2–4, 6, 7, 9, 11–26, 28, 36–39, 41, 42, 44, 46–61, and 63*

We have reviewed Petitioner’s contentions regarding dependent claims 2–4, 6, 7, 9, 11–26, and 28, as well as independent claim 36, which recites an LPHDR execution unit and the imprecision limitation, and claims 37–39, 41, 42, 44, 46–61, and 63, which depend from claim 36. *See* Pet. 36–38, 49–56. Petitioner explains how each limitation of the dependent claims is taught or rendered obvious by the disclosure of Dockser or the combination of Dockser and MacMillan. *Id.* For example, Petitioner argues that “[c]laims 11–17 each recite a minimum X and/or Y percentage higher than claim 1’s, the highest being X=10% and Y=0.2% in claim 17,” and Mr. Goodin’s software program “demonstrates that when retaining 9 mantissa fraction bits as Dockser . . . discloses, Dockser’s register bit-dropping produces” a Y relative error of at least 0.2% for 14.6% of

possible valid inputs.<sup>5</sup> *Id.* at 51–52. Petitioner’s contentions are supported by the testimony of Mr. Goodin and are persuasive based on the current record. *See id.*; Ex. 1003 ¶¶ 311–322, 361–362, 375–409. Patent Owner does not argue the challenged dependent claims separately, only disputing Petitioner’s arguments regarding the challenged independent claims and dependent claims 5, 8, 10, 40, 43, and 45. *See* Prelim. Resp. 8–28. Based on the current record, we disagree with Patent Owner’s arguments regarding the LPHDR execution unit, imprecision limitation, and limitations of the dependent claims for the reasons explained above. *See supra* Sections II.C.2, II.E.3.

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 2–4, 6, 7, 9, 11–26, 28, 36–39, 41, 42, 44, 46–61, and 63 are unpatentable over Dockser and MacMillan.

*F. Obviousness Ground Based on Dockser, Tong, and MacMillan*

*1. Claims 1–26, 28, 32–61, 63, and 67–70*

Petitioner asserts that it would have been obvious to operate the FPPs in the Dockser-MacMillan combination at Tong’s precision levels. Pet. 56–59 (citing Ex. 1008, 278, Table IV). Petitioner concludes that using Dockser’s FPP with Tong’s precision levels in MacMillan’s multiple PEs

---

<sup>5</sup> With respect to claims 11–17, we determine that Petitioner’s arguments regarding the software program are sufficient for purposes of institution. The parties are encouraged to address in their papers during trial the sufficiency of Petitioner’s algebraic analysis, premised on “retaining 9 fraction bits” for claims 11 and 12, “retaining 8 fraction bits” for claims 13–16, and “retaining 7 fraction bits” for claim 17. *See* Pet. 52–53.

would have achieved “supercomputer performance” while conserving power. *Id.* (citing Ex. 1003 ¶ 412).

Patent Owner argues that “Petitioner relies on the same flawed reasoning as discussed [in the Preliminary Response] with respect to Dockser and Tong,” but does not present separate arguments directed to the combination of Dockser, Tong, and MacMillan. Prelim. Resp. 28. We disagree with Patent Owner’s arguments based on the current record for the reasons explained above. *See supra* Sections II.C.2, II.D.2, II.E.2. On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 1–26, 28, 32–61, 63, and 67–70 are unpatentable over Dockser, Tong, and MacMillan.

*2. Petitioner’s Alternative Interpretation of Claims 5, 6, 8, 10–18, 35–61, 63, and 67–70*

Claim 5 recites that “the number of LPHDR execution units in the device exceeds by at least ten the non-negative integer number of execution units in the device *adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide*” (emphasis added). Independent claims 36 and 68, and dependent claims 8, 10, 35, 40, 43, 45, and 70, recite similar limitations, modifying the number by which the number of LPHDR execution units exceeds the number of 32-bit execution units.

Petitioner provides an “alternative interpretation” of the italicized adapted-to clause. Pet. 61. In particular, Petitioner argues that a unit could meet the adapted-to clause if it is capable of 32-bit multiplication in *some* configurations. *Id.* Under this interpretation, Dockser’s FPP would not be such a unit because it has registers and multiplier with less than 32 bits, and

thus is not capable of 32-bit multiplication in the proposed combination—rather, only the host CPU floating-point unit would meet the adapted-to clause. *Id.* (citing Ex. 1003 ¶ 427). According to Petitioner, a person of ordinary skill in the art “would have been motivated to customize Dockser’s FPPs in MacMillan’s PEs to only operate at precision levels lower than full FP 32-bit operations, in view of Tong’s teachings that ‘the fine precision of the 23-bit mantissa is not essential.’” *Id.* at 59–60.

Patent Owner argues that Petitioner does not provide a specific level of precision for its combination. Prelim. Resp. 30 (citing Pet. 60–61). Patent Owner argues that “Tong suggests 11 bits of precision for certain signal processing applications, and Petitioner provides no analysis of whether 11 bits of precision would meet the imprecision limitation.” *Id.* Also, Patent Owner argues that Petitioner does not show that, for example, one bit lower than full 32-bit operation would meet the claimed imprecision limitation. *Id.* at 32.

We disagree with Patent Owner based on the current record. In its “alternative interpretation” of claims 5, 6, 8, 10–18, and 35–70, Petitioner states that the “selected precision levels are unchanged from [the other three asserted grounds].” Pet. 60–61. In the asserted ground based on Dockser and Tong, for example, Petitioner states that Tong’s “teaching that retaining 5 mantissa fraction bits is sufficient in some applications (including ALVINN and Sphinx[]) would have motivated a [person of ordinary skill in the art] to configure Dockser’s FPP . . . to operate at a selected precision level retaining as few as 5 mantissa fraction bits.” *Id.* at 41–42. Petitioner also used a software program to determine the relative error when retaining 5 mantissa fraction bits—i.e., the number of bits that Dockser would use under Petitioner’s proposed combination with Tong. *Id.* at 42 (citing

Ex. 1003 ¶ 336). In at least this way, Petitioner has adequately explained, for the purposes of institution and on this preliminary record, which of Tong’s precision levels would be used in the proposed combination.

Patent Owner also argues that each reference teaches away from the proposed combination. Prelim. Resp. 31. In particular, Patent Owner argues that (1) Dockser teaches away because “Dockser is directed entirely to selectable precision; and disparages the use of non-selectable precision units,” (2) Tong teaches away because “Tong devotes much of its discussion to the benefits of variable or selectable precision,” and (3) MacMillan teaches away because “MacMillan is primarily concerned with providing increased performance without increased cost.” *Id.*

At this stage and on this record, we preliminarily determine that the references do not teach away from the claimed devices. The mere disclosure of more than one alternative does not constitute a teaching away. *In re Fulton*, 391 F.3d 1195, 1201 (Fed. Cir. 2004). Specifically, Patent Owner points to a statement that Tong makes about power savings: “This FP bitwidth reduction can deliver a significant power savings through the use of a variable bitwidth FP unit.” Prelim. Resp. 31 (quoting Ex. 1008, 273). Yet Tong’s statement does not criticize, discredit or otherwise discourage the claimed invention’s approach so as to teach away from it. *See Fulton*, 391 F.3d at 1201. Instead, the statement merely explains the benefits provided by Tong’s alternative approach. Ex. 1008, 273.

Likewise, although Patent Owner argues that the combination with MacMillan would increase manufacturing costs, Patent Owner does not provide sufficient evidence to support this position. *See* Prelim. Resp. 31. Even assuming MacMillan warns against the costs, Patent Owner’s argument is unsupported by any specific cost analysis. *Id.* To be sure,

Patent Owner may introduce evidence that supports this argument during trial. We also note that the fact “[t]hat a given combination would not be made by businessmen for economic reasons does not mean that persons skilled in the art would not make the combination because of some technological incompatibility,” for example. *In re Farrenkopf*, 713 F.2d 714, 718 (Fed. Cir. 1983). But, at this stage and on this record, Patent Owner’s argument is unpersuasive.

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 5, 6, 8, 10–18, 35–61, 63, and 67–70 are unpatentable over Dockser, Tong, and MacMillan.

### III. CONCLUSION

Based on the arguments presented in the Petition, we conclude that Petitioner has demonstrated a reasonable likelihood of prevailing with respect to at least one claim of the ’273 patent challenged in the Petition. Accordingly, we institute a trial on all claims and all grounds asserted in the Petition. The Board has not made a final determination under 35 U.S.C. § 318(a) with respect to the patentability of the challenged claims.

### IV. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that, pursuant to 35 U.S.C. § 314(a), an *inter partes* review of claims 1–26, 28, 32–61, 63, and 67–70 of the ’273 patent is instituted with respect to all grounds set forth in the Petition; and

FURTHER ORDERED that, pursuant to 35 U.S.C. § 314(c) and 37 C.F.R. § 42.4(b), *inter partes* review of the ’273 patent shall commence

IPR2021-00179

Patent 8,407,273 B2

on the entry date of this Decision, and notice is hereby given of the institution of a trial.

IPR2021-00179

Patent 8,407,273 B2

For PETITIONER:

Elisabeth Hunt

Richard Giunta

Anant Saraswat

WOLF, GREENFIELD & SACKS, P.C.

ehunt-ptab@wolfgreenfield.com

rgiunta-ptab@wolfgreenfield.com

asaraswat-ptab@wolfgreenfield.com

For PATENT OWNER:

Peter Lambrianakos

Vincent Rubino

Enrique Iturralde

FABRICANT LLP

plambrianakos@fabricantllp.com

vrubino@fabricantllp.com

eiturralde@fabricantllp.com